

Claims

We claim:

1. A method for allocating system resources among groups having entitlement values and maximum limits comprising:

5 allocating a computer system resource to active groups according to respective entitlement values;

determining an excess entitlement allocated to inactive groups; and

10 reallocating the excess entitlement to the active groups in proportion to the respective entitlement values, without exceeding a maximum limit for each of the active groups,

whereby each of the active groups has a maximal value representing the groups' total proportionate share of the system resource after the excess entitlement has been reallocated.

15 2. The method of claim 1, wherein the maximal values for inactive groups is set equal to zero.

3. The method of claim 1, further comprising:
calculating a scaling ratio for each group; and
sorting active groups by their scaling ratios.

20 4. The method of claim 3, wherein the scaling ratio is a ratio between the maximum limit and the entitlement value.

5. The method of claim 3, wherein the step of reallocating comprises:
determining whether unprocessed groups can scale by the scaling ratio of a current group without exhausting unallocated resources; and
25 if the unprocessed groups can scale without exhausting the unallocated resources, then setting the maximal value of the current group equal to the maximum limit of the current group.

6. The method of claim 5, wherein the step of reallocating further comprises:
if the unprocessed groups cannot scale without exhausting the unallocated resources, then scaling the unprocessed groups by the unallocated resources.

7. The method of claim 3, further comprising processing the groups individually as sorted by the scaling ratios, whereby the groups having a higher maximum limit relative to their entitlement values are processed after groups having a lower maximum limit relative to their entitlement values.

8. A computer system for allocating system resources among groups having entitlement values and maximum limits comprising:

a memory; and

a software program stored in the memory, which software program

allocates a system resource to active groups according to respective entitlement values;

determines an excess entitlement allocated to inactive groups; and

reallocates the excess entitlement to the active groups in proportion to the respective entitlement values, without exceeding a maximum limit for each of the active groups,

whereby each active group has a maximal value representing the groups' total proportionate share of the system resource after the excess entitlement has been reallocated.

9. The computer system of claim 8, wherein the software program calculates a scaling ratio for each group; and sorts active groups by their scaling ratios.

10. The computer system of claim 9, wherein the scaling ratio is a ratio between the maximum limit and the entitlement value.

11. The computer system of claim 9, wherein the step of reallocating comprises:

determining whether unprocessed groups can scale by the scaling ratio of a current group without exhausting unallocated resources; and

if the unprocessed groups can scale without exhausting the unallocated resources, then setting the maximal value of the current group equal to the maximum limit of the current group.

12. The method of claim 11, wherein the step of reallocating further comprises:

if the unprocessed groups cannot scale without exhausting the unallocated resources, then scaling the unprocessed groups by the unallocated resources.

13. The computer system of claim 12, wherein the software program processes the groups individually as sorted by the scaling ratios, whereby the groups having a higher maximum limit relative to their entitlement values are processed after groups having a lower maximum limit relative to their entitlement values.

14. The computer system of claim 13, wherein the step of reallocating further comprises, if a portion of the excess entitlement remains unallocated after processing all active groups, reallocating the portion to one or more active or inactive groups.

15. A software system comprising:
a tangible storage medium; and
a software program stored in the medium, which software program
allocates a system resource to active groups according to respective entitlement values;
determines an excess entitlement allocated to inactive groups; and
reallocates the excess entitlement to the active groups in proportion to active groups' respective entitlement values, without exceeding a maximum limit for the groups,

whereby each group has a maximal value representing the groups' total proportionate shares of the resource after the excess entitlement has been reallocated.

16. The computer system of claim 15, wherein the software program calculates a scaling ratio for each group; and sorts active groups by their scaling ratios.

17. The computer system of claim 16, wherein the scaling ratio is a ratio between the maximum limit and the entitlement value.

18. The computer system of claim 16, wherein the step of reallocating comprises:
determining whether unprocessed groups can scale by the scaling ratio of a current group without exhausting unallocated resources; and
if the unprocessed groups can scale without exhausting the unallocated resources, then setting the maximal value of the current group equal to the maximum limit of the current group.

19. The method of claim 18, wherein the step of reallocating further comprises:

if the unprocessed groups cannot scale without exhausting the unallocated resources, then scaling the unprocessed groups by the unallocated resources.

20. The computer system of claim 19, wherein the software program processes the groups individually as sorted by the scaling ratios, whereby the groups having a higher maximum limit relative to their entitlement values are processed after groups having a lower maximum limit relative to their entitlement values.

5